

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

APPLICANT:	CASTELLI, et al.	DOCKET NO:	YOR9-2001-0461 (8728-524)
SERIAL NO:	09/916,935	ART UNIT:	2121
FILED:	July 27, 2001	EXAMINER:	HIRL, Joseph P.
FOR:	METHOD AND APPARATUS FOR PREDICTION OF COMPUTER SYSTEM PERFORMANCE BASED ON TYPES AND NUMBERS OF ACTIVE DEVICES		

Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

In response to the Final Office Action dated February 1, 2006, rejecting Claims 1-40 under 35 U.S.C. §102(e) and Claims 42 and 7 under 35 U.S.C. § 112, first and second paragraphs, respectively, Applicant appeals pursuant to the Notice of Appeal dated May 1, 2006, and submits this appeal brief.

TABLE OF CONTENTS

	Page
1. REAL PARTY IN INTEREST	1
2. RELATED APPEALS AND INTERFERENCES	1
3. STATUS OF CLAIMS	1
4. STATUS OF AMENDMENTS	1
5. SUMMARY OF CLAIMED SUBJECT MATTER	1
6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL	9
7. ARGUMENTS	9
I. Claim 42 Complies With The Written Description Requirement of 35 U.S.C. § 112, First Paragraph, and Does Not Contain New Subject Matter	9
A. The Subject Matter of Claim 42 Is Contained Within the Specification	10
B. The Deemed “New Subject Matter” Is Not Contained Within Claim 42	11
II. Claim 7 Complies With The Written Description Requirement of 35 U.S.C. § 112, Second Paragraph, and the Term “Rejuvenating” Is Not Indefinite	11
A. The Term “Rejuvenating” Is Not a “Relative Term”	11
B. The Intent of a Claim Is Not Relevant	12
C. “Monitoring, over a period of time” Is Not Perpetual	13
D. The Term “Rejuvenating” Is Previously Known	14
III. The Teachings of Sweet, et al. Do Not Support the Anticipation Rejections	15
A. The Subject Matter Claimed In Each of Claims 1-40, Taken As a Whole, Is Not Anticipated By Sweet, et al.	16
1. Sweet, et al. Does Not Anticipate Claim 1	17

2. Sweet, et al. Does Not Anticipate Claim 21	18
3. Sweet, et al. Does Not Anticipate Claim 31	19
4. Sweet, et al. Does Not Anticipate Claim 38	20
5. Sweet, et al. Does Not Anticipate Claim 39	21
6. Sweet, et al. Does Not Anticipate Claim 40	22
7. Sweet, et al. Does Not Anticipate Claims 2-20, 22-30 and 32-37	23
IV. Conclusion	23
CLAIMS APPENDIX	24
EVIDENCE APPENDIX	35
RELATED PROCEEDINGS APPENDIX	36

1. REAL PARTY IN INTEREST

The real party in interest is International Business Machines Corporation, by virtue of an assignment dated January 15, 2002 and January 24, 2002, and recorded February 20, 2002 in the U.S. Patent and Trademark Office at reel 012640 and frame 0264.

2. RELATED APPEALS AND INTERFERENCES

There are no related commonly assigned U.S. patent applications under appeal.

3. STATUS OF CLAIMS

Claims 1-40 and 42 are pending and stand rejected. Claims 1-40 and 42 are under appeal. A copy of the claims under appeal is presented in the Claims Appendix.

4. STATUS OF AMENDMENTS

No amendments have been filed after the Final Office Action in this case.

5. SUMMARY OF CLAIMED SUBJECT MATTER

For purposes of illustration, the invention of claims 1, 7, 21, 31, 38, 39, 40, and 42 will be described with reference to the exemplary figures and the corresponding text of Applicant's specification. It is to be understood that the following description of the claimed subject matter and references to the specification and drawings are for illustrative purposes to provide some context for the claimed subject matter, but nothing herein shall be construed as placing any limitations thereon.

In general, the claimed subject matter provides an apparatus and method to predict resource utilization in a computer system having a plurality of devices that includes monitoring, resource utilization and the number of active devices to obtain monitored values of the resource utilization and the number of active devices. The subsequent resource utilization is predicted, based upon the monitored values of the resource utilization and the number of active devices. Additional methods identify resource saturation and predict the effects of adding a new device in the computer system.

For instance, FIG. 1 of the specification illustrates a computer system (101) monitored by a capacity management system (150) that includes a data collection module (102), a user interface (103), a monitoring and displaying facility (104), a preprocessing module (105), and a forecasting module (106). The forecasting module (106) outputs a prediction (107).

Monitoring and prediction software acquires data from a plurality of devices including the number of active devices (i.e., those device that are turned on) at each point in time. Prediction algorithms illustrated in FIGs. 2-8 and 10-17 are used in exemplary embodiments to implement resource utilization predictions.

Claim 1 recites:

*A method for predicting a subsequent resource utilization in a computer system having a plurality of devices, comprising the steps of:
monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices; and
predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.*

An exemplary embodiment of claim 1 may be described with reference to FIGs. 1, 2-8, and 13, and to relevant descriptions in specification page 1, lines 10-13; page 10, lines 10-12; page 10, line 17 to page 12, line 7; page 13, line 6 to page 22, line 4; and page 26, lines 7-22.

A method for predicting a subsequent resource utilization (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26, lines 7-22) in a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), comprising the steps of:

monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and

predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26, lines 7-22).

Claim 7 recites:

The method of Claim 1, wherein the computer system comprises computer software, and said method further comprising the step of rejuvenating the computer software, based upon the predicted subsequent resource utilization.

An exemplary embodiment of claim 7 may be described with reference to FIG. 1, to relevant descriptions in specification page 2, lines 7-16; page 13, lines 3-5, and to commonly assigned U.S. Patent Application Ser. No. 09/706,737, filed on Nov. 7, 2000, now issued as U.S. Patent No. 6,993,458, entitled "Method and Apparatus for Preprocessing Technique for Forecasting in Capacity Management, Software Rejuvenation and Dynamic Resource Allocation Applications," the disclosure of which was incorporated by reference in the present application. (page 2, lines 7-16). It is well known in the art that computer systems comprise computer software.

The method of Claim 1, wherein the computer system (FIG. 1, 101, page 10, lines 10-12) comprises computer software (known in the art), and said method further comprising the step of

rejuvenating the computer software, based upon the predicted subsequent resource utilization
(page 2, lines 7-16, page 13, lines 3-5).

Claim 21 recites:

A method for identifying resource saturation in a computer system having a plurality of devices, comprising the steps of:

monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.

An exemplary embodiment of claim 21 may be described with reference to FIGs. 1 and 9-12, and to relevant descriptions in specification page 1, lines 10-13; page 10, lines 10-12; page 10, line 17 to page 12, line 7; page 13, lines 6-16; and page 22, line 10 to page 26, line 6.

A method for identifying resource saturation (FIGs. 9-12, page 22, line 10 to page 26, line 6) in a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), comprising the steps of:

monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and

identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices (FIGs. 9-12, page 22, line 10 to page 26, line 6).

Claim 31 recites:

A method for predicting effects of adding a new device on a computer system having a plurality of devices, the method comprising the steps of:

monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices.

An exemplary embodiment of claim 31 may be described with reference to FIGs. 1 and 16-17, and to relevant descriptions in specification page 1, lines 10-13; page 10, lines 10-12; page 10, line 17 to page 12, line 7; page 13, lines 6-16; and page 30, line 4 to page 31, line 22.

A method for predicting effects of adding a new device (FIGs. 16-17, page 30, line 4 to page 31, line 22) on a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), the method comprising the steps of:

monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and

predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices (FIGs. 16-17, page 30, line 4 to page 31, line 22).

Claim 38 recites:

An apparatus for predicting a subsequent resource utilization in a computer system having a plurality of devices, comprising:

a monitoring device for monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices; and

a prediction device for predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.

An exemplary embodiment of claim 38 may be described with reference to FIGs. 1, 2-8, and 13, and to relevant descriptions in specification page 1, lines 10-13; page 10, line 10 to page 22, line 4; and page 26, lines 7-22.

An apparatus (FIG. 1, 150, page 10, line 10 to page 13, line 5) for predicting a subsequent resource utilization (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26,

lines 7-22) in a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), comprising:

a monitoring device (FIG. 1 102, page 11, lines 3-5) for monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and

a prediction device (FIG. 1, 106, page 12, lines 17-19) for predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26, lines 7-22).

Claim 39 recites:

An apparatus for identifying resource saturation in a computer system having a plurality of devices, comprising:

a monitoring device for monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

a forecasting device for identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.

An exemplary embodiment of claim 39 may be described with reference to FIGs. 1 and 9-12, and to relevant descriptions in specification page 1, lines 10-13; page 10, line 10 to page 13, line 16; and page 22, line 10 to page 26, line 6.

An apparatus (FIG. 1, 150, page 10, line 10 to page 13, line 5) for identifying resource saturation (FIGs. 9-12, page 22, line 10 to page 26, line 6) in a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), comprising:

a monitoring device (FIG. 1, 102, page 11, lines 3-5) for monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the

resource utilization and the number of active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and

a forecasting device (FIG. 1, 106, page 12, lines 17-19) for identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices (FIGs. 9-12, page 22, line 10 to page 26, line 6).

Claim 40 recites:

An apparatus for predicting effects of adding a new device on a computer system having a plurality of devices, the apparatus comprising:

a monitoring device for monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

a forecasting device for predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices.

An exemplary embodiment of claim 40 may be described with reference to FIGs. 1 and 16-17, and to relevant descriptions in specification page 1, lines 10-13; page 10, line 10 to page 13, line 16; and page 30, lines 4 to page 31, line 22.

An apparatus (FIG. 1, 150, page 10, line 10 to page 13, line 5) for predicting effects of adding a new device (FIGs. 16-17, page 30, line 4 to page 31, line 22) on a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), the apparatus comprising:

a monitoring device (FIG. 1, 102 , page 11, lines 3-5) for monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and

a forecasting device (FIG. 1, 106, page 12, lines 17-19) for predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices (FIGs. 16-17, page 30, lines 4 to page 31, line 22).

Claim 42 recites:

A method for predicting a subsequent resource utilization in a computer system having a plurality of devices, the plurality of devices comprising active devices and non-active devices, comprising the steps of:

monitoring, over a period of time, a contemporaneous resource utilization to obtain first monitored values of the contemporaneous resource utilization;

monitoring, over the period of time, a number of the active devices to obtain second monitored values of the number of the active devices, wherein the monitored number is capable of varying over the period of time;

monitoring, over the period of time, a type of each of the active devices to obtain third monitored values of the type of the each of the active devices; and

predicting the subsequent resource utilization, based upon the first monitored values, the second monitored values, and the third monitored values.

An exemplary embodiment of claim 42 may be described with reference to FIGs. 1-8 and 13-15, and relevant descriptions of specification page 1, lines 10-13; page 10, lines 10-12; page 10, line 17 to page 12, line 7; page 13, line 6 to page 22, line 4; page 25, lines 13-20; page 26, lines 7-22; page 27, line 16 to page 28, line 15; and page 29, lines 7-13.

A method for predicting a subsequent resource utilization (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26, lines 7-22) in a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), the plurality of devices comprising active devices and non-active devices (FIG. 14, page 27, line 16 to page 28, line 15), comprising the steps of:

monitoring, over a period of time, a contemporaneous resource utilization to obtain first monitored values of the contemporaneous resource utilization (page 10, line 17 to page 12, line 7, page 13, lines 6-16);

monitoring, over the period of time, a number of the active devices (page 10, lines 17-22) to obtain second monitored values of the number of the active devices (page 13, lines 6-16), wherein the monitored number is capable of varying over the period of time (page 25, lines 13-20);

monitoring, over the period of time, a type of each of the active devices to obtain third monitored values of the type of the each of the active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and

predicting the subsequent resource utilization, based upon the first monitored values, the second monitored values, and the third monitored values (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26, lines 7-22, FIG. 15, 1501, page 29, lines 7-13).

6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

- I. Claim 42 was rejected under 35 U.S.C. § 112, first paragraph, for containing subject matter not described in the specification.
- II. Claim 7 was rejected under 35 U.S.C. § 112, second paragraph, because the term “rejuvenating” was deemed a relative term rendering the claim indefinite.
- III. Claims 1-40 were rejected under 35 U.S.C. § 102(e), as being anticipated by Sweet, et al.

7. ARGUMENTS

- I. **Claim 42 Complies With the Written Description Requirement of 35 U.S.C. § 112, First Paragraph, and Does Not Contain New Subject Matter**

The following is a quotation of the first paragraph of 35 U.S.C. 112:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

A. The Subject Matter of Claim 42 Is Contained Within The Specification

All of the subject matter of claim 42 was present in the specification as originally filed.

Claim 42 recites, “A method for predicting a subsequent resource utilization (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26, lines 7-22) in a computer system (FIG. 1, 101, page 10, lines 10-12) having a plurality of devices (page 1, lines 10-13), the plurality of devices comprising active devices and non-active devices (FIG. 14, page 27, line 16 to page 28, line 15), comprising the steps of: monitoring, over a period of time, a contemporaneous resource utilization to obtain first monitored values of the contemporaneous resource utilization (page 10, line 17 to page 12, line 7, page 13, lines 6-16); monitoring, over the period of time, a number of the active devices (page 10, lines 17-22) to obtain second monitored values of the number of the active devices (page 13, lines 6-16), wherein the monitored number is capable of varying over the period of time (page 25, lines 13-20); monitoring, over the period of time, a type of each of the active devices to obtain third monitored values of the type of the each of the active devices (page 10, line 17 to page 12, line 7, page 13, lines 6-16); and predicting the subsequent resource utilization, based upon the first monitored values, the second monitored values, and the third monitored values (FIGs. 2-8, page 13, line 6 to page 22, line 4, FIG. 13, page 26, lines 7-22, FIG. 15, 1501, page 29, lines 7-13).”

B. The Deemed “New Subject Matter” Is Not Contained Within Claim 42

The basis for the Examiner’s rejection is that the specification and drawings do not contain “segmenting of variable length historic data, comparison of segmented historic data, comparison of unequal length of historic data and comparison of non-homogeneous type data.” As can be clearly seen from the text of claim 42 above, the subject matter the Examiner deems “new matter,” such as “segmented historic data,” is not recited in claim 42.

Thus, Applicant respectfully submits there is no new matter in claim 42 and believes there is no applicable rejection under 35 U.S.C. § 112, first paragraph, to claim 42. Accordingly, since there have been no other grounds of rejection regarding claim 42, Applicant respectfully requests that the rejection be withdrawn and that claim 42 be allowed.

II. Claim 7 Complies With The Written Description Requirement of 35 U.S.C. § 112, Second Paragraph, and the Term “Rejuvenating” Is Not Indefinite

A. The Term “Rejuvenating” Is Not a “Relative Term”

MPEP § 2173.05(b) relates to “relative terminology,” which pertains to terms of degree and variable objects. The Examiner has incorrectly invoked this section of the MPEP by rejecting claim 7 under 35 U.S.C. § 112, second paragraph, because he deemed the term “rejuvenating” a “relative term” rendering the claim indefinite. Additionally, the Examiner did not provide a full explanation of the deficiency as required by MPEP 706.03(d). The Examiner merely stated that the term “rejuvenating” was a “relative term,” without giving any support for the rejection, such as in what way the term was “relative.” This rejection and its explanations are fundamentally flawed for several reasons.

A term of degree is a modifying word, for instance a word modifying an object by describing its extent, such as a “tall” or “heavy” object. A variable object is an object that can have indeterminate characteristics, such as “the height of the rider that the bicycle was designed for.” “Rejuvenating” describes a process used as a method step and is not a term of degree or a variable object. “Rejuvenating” does not describe the extent of something and is not a reference to an object that is variable.

B. The Intent of a Claim Is Not Relevant

In a response to Applicant’s arguments, the Examiner still did not give any support for his rejection that the term rejuvenating was “relative.” Instead, the Examiner responded with a different alleged deficiency unrelated to relative terminology. Specifically, the Examiner stated, “Claim 1. teaches ‘predicting subsequent resource utilization in a computer system’ based on monitoring resources over time which requires computer system operation to realize the limitations. Rejuvenation or restarting of the software is a discontinuous operation and teaches away from the intent of claim 1 wherein resources are monitored over time ... a shutdown monitors nothing. Restarting means that all resource utilization returns to zero ... stops. The intent of claim 1. is for continuous operation as indicated by ‘monitoring over a period of time, a contemporaneous resource utilization.’ The limitations of claim 7 are uncertain or indefinite to claim 1.” Nowhere did the Examiner address anything related to the term “rejuvenating” being a relative term. However, the Examiner described new issues related to “teaching away” from the “intent” of a claim and the limitations of one claim being uncertain or indefinite to another claim. The new issues regarding teaching away from the intent of an element of a claim from which this claim depends are not deficiencies Applicant can find in the patent laws, regulations or MPEP.

The "intent" of a claim has no legal significance. The "intent" of a claim, even if significant, could not be defined by one of the claim elements. Whether or not claim 7 has an "intent" of a continuous operation, there is no "intent" for a perpetual operation. It is irrelevant whether the monitoring operation is continuous or discontinuous. A "period of time" has an end, occurring, for example, when a condition is met. Claim 7 describes an operation to be performed following the occurrence of a condition. Claim 7 is a further limitation to claim 1 and not necessarily a further limitation to the monitoring step of claim 1 as implied by the Examiner.

C. "Monitoring, over a period of time" Is Not Perpetual

"Monitoring, over a period of time" is not a perpetual operation but is performed for the period of time. The method of claim 1 provides method steps such that monitoring yields obtained results that prediction is based upon. In the case of claim 7, if the obtained results warrant, an additional step of rejuvenating the software occurs. It is irrelevant whether the monitoring operation is continuous or discontinuous. At some point in time, other method steps occur, whether or not monitoring continues to occur. Monitoring may be carried out by a system other than the computer system or software that is to be monitored and rejuvenated.

Also, the Examiner appears to be reading into claim 1 limitations that are not recited in claim 1 or limitations in the specification that would limit the scope of claim 1. The Examiner states that the word "monitoring" in one of the claim elements implies that the "intent" of claim 1 is that it be a continuous operation and that the method step of rejuvenating the software in claim 7 is a discontinuous operation and "teaches away" from the "intent" of claim 1.

Additionally, the Examiner has equated rejuvenation with restarting and then restarting with system shutdown. Applicant does not disclose the discontinuous operation of system

shutdown as Examiner has characterized rejuvenating. Applicant's disclosure does not use the word "shutdown" to describe rejuvenation.

Furthermore, in the Advisory Action the Examiner asks how much rejuvenation is to be performed. Software rejuvenation is a discrete operation and not an operation that can be performed by percentages as indicated by the Examiner.

D. The Term "Rejuvenating" Is Previously Known

Software rejuvenation is described in commonly assigned U.S. Patent Application Ser. No. 09/706,737, filed Nov. 7, 2000, now U.S. Patent No. 6,993,458, entitled "Method and Apparatus for Preprocessing Technique for Forecasting in Capacity Management, Software Rejuvenation and Dynamic Resource Allocation Applications," the disclosure of which was incorporated by reference in the present application. Software rejuvenation, as recited in claim 7 and disclosed in the above reference, is described below.

Software rejuvenation is a discipline concerned with scheduling the termination and re-initialization of applications or operating systems, in order to avoid catastrophic crashes. It is the mechanical equivalent of preventive maintenance. The assumption is that "bugs" in software might cause programs to allocate resources and never release them. Eventually, a crash is caused when the needed resource is exhausted. A typical example of this kind of "bug" is a memory leak, where an application allocates memory but mistakenly never releases it. Other resources that can be exhausted are semaphores, mutexes, handles etc.

The motivation for rejuvenating software rather than waiting for a crash is twofold. A first reason is to prevent data loss and secondly to guarantee quality of service. For example, a crash of a database usually requires a rollback operation to a consistent state (reliably stored on persistent media during periodic checkpoint operations), the reconstruction of all the memory-based data structures, and the re-execution of all the transactions that have been committed after the most recent checkpoint, and have been stored in an appropriate log. The time required to recover from a catastrophic crash can be in the order of hours to tens of hours, during which the database is not available to process new transactions. Since the re-initialization of the database to a checkpointed state is

the least expensive of the above operations, it is beneficial to schedule the rejuvenation of the piece of software right after a checkpoint operation, possibly during a time of low utilization.

In its simplest form, rejuvenation is based on static scheduling, where the application or the operating system is periodically restarted. ... If a resource is in danger of being exhausted in the near future, then the system administrator is notified of the problem and decides whether a rejuvenation should be appropriately scheduled.

... Then, using the prediction techniques, the software package can estimate the probability that a crash happens if a scheduled rejuvenation is not executed. If this probability is acceptably low, then the rejuvenation step is skipped.

As disclosed, rejuvenation does not require a shutdown of the computer system. The computer software may be restarted and re-initialized. Additionally, it is irrelevant whether there is a shutdown of the computer system since the method is not necessarily implemented by the same computer system or software as is monitored and rejuvenated.

Accordingly, claim 7 satisfies the requirements under 35 U.S.C § 112, second paragraph and Applicant respectfully requests withdrawal of the rejection.

III. The Teachings of Sweet, et al. Do Not Support the Anticipation Rejections

For a claim to be anticipated under 35 U.S.C. § 102, all elements of the claim must be found in a single prior art reference (see, e.g., *Scripps Clinic & Research Found. V. Genentech, Inc.*, 927 F.2d 1565, 1576, 18 U.S.P.Q.2d 1001, 1010 (Fed. Cir. 1991)). The identical invention must be shown in as complete detail as is contained in the claim. (See MPEP § 2131). The single prior art reference must disclose all of the elements of the claimed invention functioning essentially in the same manner (see, e.g., *Shanklin Corp. v. Springfield Photo Mount Corp.*, 521 F.2d 609 (1st Cir. 1975)).

Here, Applicant respectfully asserts that Sweet is legally deficient to establish a *prima facie* case of anticipation against any of claims 1-40. At the very least, Sweet does not anticipate independent claims 1, 21, 31, 38, 39 and 40, for the following reasons.

A. The Subject Matter Claimed In Each of Claims 1-40, Taken As a Whole, Is Not Anticipated By Sweet, et al.

Claims 1-40 were rejected under 35 U.S.C. § 102(e), as being anticipated by Sweet, et al., U.S. Patent No. 6,836,800 (hereinafter, Sweet). Sweet is concerned with resource management related to network traffic loads and anticipating network slow-downs, such as from software application response time. Nowhere is Sweet concerned with the addition of new devices or resources, such as Applicant's "subsequent resource utilization," or their prediction and effects.

Independent claims 1, 38, and 42 have limitations including "subsequent resource utilization" and Independent claims 31 and 40 have limitations including "adding a new device," which are directly related to new devices or resources on a computer system. Since Sweet does not disclose anything related to the addition of new devices or to subsequent resource utilization. Therefore, Sweet does not anticipate these claim limitations of independent claims 1, 31, 38, 40 and 42.

Further, Sweet discloses monitoring a number of network devices, computing statistics and displaying statistical information for selected key network devices. However, Sweet does not disclose "identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices," as recited in Applicant's independent claims 21 and 39. Therefore, Sweet does not anticipate independent claims 21 and 39.

In order for a reference to anticipate a claim, each and every element set forth in the claim must be found, either expressly or inherently, in the reference. Verdegaal Bros. V. Union Oil Co. of California, 814 F.2d 628,631, 2 USPQ2d 1051,1053 (Fed. Cir. 1987). The reference cited by the Examiner does not disclose, either expressly or inherently, each and every element of Applicant's Claims 1, 21, 31, 38, 39, 40, or 42. Therefore, the rejections under 35 U.S.C. § 102(e) should be reversed.

1. Sweet, et al. Does Not Anticipate Claim 1

The Examiner states that Sweet anticipates, *inter alia*, “predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices (Sweet, c 3, l 25-41; Examiner's Note (EN): para 11 applies; the number of active devices is integrated into the system operation as shown in Fig. 1).”

Sweet discloses that “data is gathered about the network (step 1010), the gathered data is analyzed to determine whether a signature exists (step 1020), and if a signature exists, the signature is used for purposes such as generating alarms for unusual activity (or inactivity), reporting on the status of the network, and planning changes such as upgrades to the network (step 1030).” Sweet further discloses that “The signature detection software ... derives signature data 40 and other data 42 from the gathered data. The other data may include data needed for ... trending predictions.” Sweet suggests that “trending” relates to “the capacity or configuration of the network.”

However, Sweet does not disclose “predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.” Applicant respectfully submits that analyzing data, generating alarms, reporting status

and planning changes do not constitute or suggest “predicting the subsequent resource utilization.” Furthermore, Applicant respectfully submits that trending predictions as disclosed by Sweet are predictions of network performance about the capacity or configuration of the network, and not “predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.” Sweet indicates that the “other data” that may include data needed for trending predictions is derived from signature detection software and not from “contemporaneous resource utilization and the number of active devices.” Accordingly, Claim 1 is neither anticipated nor rendered obvious in light of Sweet and is allowable over the prior art of record.

2. Sweet, et al. Does Not Anticipate Claim 21

The Examiner states that Sweet anticipates, *inter alia*, “identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices (Sweet, c 4, 1 42-58; EN: signature detection establishes monitored values for the active devices).” Furthermore, in the Examiner’s response to Applicant’s November 4, 2005 arguments regarding claim 21, the Examiner states that “‘identifying resource saturation’ is equivalent to an ‘alarm threshold’ which Sweet teaches @ c 2:66.’

Sweet discloses “The signature detection software includes statistical analysis software 38 that derives signature data 40 and other data 42 from the gathered data. The other data may include data needed for subsequent calculations, or data for statistical correlation or for trending predictions.” In the cited passage of the Examiner’s response, Sweet discloses “...the signature may be used to establish an alarm threshold (i.e., an alert threshold) to allow a network manager

to be alerted automatically to unusually high data traffic, due perhaps to a network malfunction or unauthorized use of the network...”

However, Sweet does not disclose “identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.” Nor does Sweet disclose that the alarm threshold is related to “resource utilization and the number of active devices.” “High data traffic,” “network malfunction” or “unauthorized use” do not equate with “resource utilization and the number of active devices.” Accordingly, claim 21 is neither anticipated nor rendered obvious in light of Sweet and is allowable over the prior art of record.

3. Sweet, et al. Does Not Anticipate Claim 31

The Examiner states that Sweet anticipates, *inter alia*, “predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices (Sweet, c 2, l 5-20; EN: such is automatic adaptation).”

Sweet discloses “... resource usage can be tracked ... enabling highly meaningful analysis and presentation of information. In the case of a network, performance thresholds can be automatically adapted and kept current, ... Rich details of network traffic patterns can be exposed and alert and alarm thresholds can be automatically tuned, allowing effective bandwidth management, capacity planning, and development of realistic service level expectations based on objective information.”

However, Sweet does not disclose “predicting the effects of adding the new device with respect to the selected monitored resource, based upon the monitored values of the resource utilization and the number of active devices.” The Examiner equates Sweet’s “performance thresholds can be automatically adapted” with Applicant’s “predicting the effects of adding the

new device.” Applicant respectfully disagrees with the Examiner’s apparent belief that “automatically adapting” is synonymous with “predicting the effects of adding the new device.” In Sweet, automatically adapting refers to performance thresholds and not to “new devices.” Accordingly, claim 31 is neither anticipated nor rendered obvious in light of Sweet and is allowable over the prior art of record.

4. **Sweet, et al. Does Not Anticipate Claim 38**

The Examiner states that Sweet anticipates, *inter alia*, “predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices (Sweet, c 2, l 5-20; EN: such is automatic adaptation).”

Sweet discloses “... resource usage can be tracked ... enabling highly meaningful analysis and presentation of information. In the case of a network, performance thresholds can be automatically adapted and kept current, ... Rich details of network traffic patterns can be exposed and alert and alarm thresholds can be automatically tuned, allowing effective bandwidth management, capacity planning, and development of realistic service level expectations based on objective information.”

However, Sweet does not disclose “predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.” The Examiner equates Sweet’s “performance thresholds can be automatically adapted” with Applicant’s “predicting the subsequent resource utilization.” Applicant respectfully disagrees with the Examiner’s apparent belief that “automatically adapting” is synonymous with “predicting the subsequent resource utilization.” In Sweet, automatically adapting refers to performance thresholds and not to “subsequent resource utilization.”

Additionally, the Examiner rejects claim 38 stating that Sweet discloses, *inter alia*, “predicting the effects of adding the new device,” which does not apply to claim 38 since the limitation “predicting the effects of adding the new device” is not contained in claim 38. Accordingly, claim 38 is neither anticipated nor rendered obvious in light of Sweet and is allowable over the prior art of record.

5. Sweet, et al. Does Not Anticipate Claim 39

The Examiner states that Sweet anticipates, *inter alia*, “predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices (Sweet, c 2, l 5-20; EN: such is automatic adaptation).”

Sweet discloses “... resource usage can be tracked ... enabling highly meaningful analysis and presentation of information. In the case of a network, performance thresholds can be automatically adapted and kept current, ... Rich details of network traffic patterns can be exposed and alert and alarm thresholds can be automatically tuned, allowing effective bandwidth management, capacity planning, and development of realistic service level expectations based on objective information.”

However, Sweet does not disclose “a forecasting device for identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.” The Examiner equates Sweet’s “performance thresholds can be automatically adapted” with Applicant’s “forecasting device for identifying resource saturation.” Applicant respectfully disagrees with the Examiner’s apparent belief that “automatically adapting” is synonymous with “forecasting device for identifying resource saturation.” In Sweet, automatically adapting refers to performance thresholds and not to “forecasting device for identifying resource saturation.”

Additionally, the Examiner rejects claim 39 stating that Sweet discloses, *inter alia*, “predicting the effects of adding the new device,” which does not apply to claim 39 since the limitation “predicting the effects of adding the new device” is not contained in claim 39. Accordingly, claim 39 is neither anticipated nor rendered obvious in light of Sweet and is allowable over the prior art of record.

6. Sweet, et al. Does Not Anticipate Claim 40

The Examiner states that Sweet anticipates, *inter alia*, “predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices (Sweet, c 2, l 5-20; EN: such is automatic adaptation).”

Sweet discloses “... resource usage can be tracked ... enabling highly meaningful analysis and presentation of information. In the case of a network, performance thresholds can be automatically adapted and kept current, ... Rich details of network traffic patterns can be exposed and alert and alarm thresholds can be automatically tuned, allowing effective bandwidth management, capacity planning, and development of realistic service level expectations based on objective information.”

However, Sweet does not disclose “a forecasting device for predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices.” The Examiner equates Sweet’s “performance thresholds can be automatically adapted” with Applicant’s “predicting the effects of adding the new device.” Applicant respectfully disagrees with the Examiner’s apparent belief that “automatically adapting” is synonymous with “predicting the effects of adding the new device.” Automatically adapting refers to performance thresholds and not to “new devices.” Accordingly, claim 40 is

neither anticipated nor rendered obvious in light of Sweet and is allowable over the prior art of record.

7. Sweet, et al. Does Not Anticipate Claims 2-20, 22-30 and 32-37

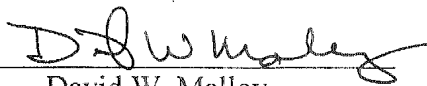
Since claims 2-20, 22-30 and 32-37 depend from independent claims 1, 21 and 31, respectively, they are allowable for at least the reasons given above for the independent claims.

IV. Conclusion

Accordingly, for at least the reasons set forth above, claims 1-40 and 42 are patentable.

Therefore, it is respectfully requested that the Board reverse all claim rejections under 35 U.S.C. § 102(e) and § 112, first and second paragraphs.

Respectfully submitted,


By: David W. Malloy
Reg. No. 56,652
Attorney for Applicants

F. CHAU & ASSOCIATES, LLC
130 Woodbury Road
Woodbury, New York 11797
Tel: (516) 692-8888
Fax: (516) 692-8889

FVC:DWM:pd

CLAIMS APPENDIX

1. A method for predicting a subsequent resource utilization in a computer system having a plurality of devices, comprising the steps of:

monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices; and

predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.

2. The method of Claim 1, wherein said predicting step further comprises the steps of:

computing a regression model of prediction parameters on the number of active devices; constructing an empirical distribution of the number of active devices; and combining the regression model and the empirical distribution to produce a prediction model.

3. The method of Claim 2, wherein said step of combining the regression model and the empirical distribution comprises the step of computing, with respect to the empirical distribution, an expected value of each of one or more of the prediction parameters.

4. The method of Claim 3, wherein said combining step further comprises the steps of:

for each of the one or more prediction parameters,

for each of the monitored values of the number of active devices, computing confidence intervals for the one or more prediction parameters; and

selecting a corresponding one of the confidence intervals for the expected value of each of the one or more prediction parameters.

5. The method of Claim 3, wherein said combining step further comprises the steps of:

computing confidence intervals for the prediction parameters, for each of the monitored values of the number of active devices; and

combining the confidence intervals with the empirical distribution, to produce modified confidence intervals for the prediction parameters.

6. The method of Claim 1, further comprising the step of managing a resource capacity of the computer system, based upon the predicted subsequent resource utilization

7. The method of Claim 1, wherein the computer system comprises computer software, and said method further comprising the step of rejuvenating the computer software, based upon the predicted subsequent resource utilization.

8. The method of Claim 1, further comprising the step of dynamically allocating at least one resource of the computer system, based upon the predicted subsequent resource utilization.

9. The method of Claim 1, wherein said predicting step comprises the steps of:
identifying any of the plurality of devices that are relevant to a monitored resource; and
restricting at least one subsequent operation of the computer system that corresponds to
the monitored resource to use only devices identified as relevant to the monitored resource from
among the plurality of devices.

10. The method of Claim 9, wherein said identifying step comprises the steps of:
for a given device currently being evaluated for relevance to the monitored resource,
computing a first regression model for the monitored resource on the number of active
devices, based upon data acquired when the given device is active;
computing a second regression model for the monitored resource on the number of active
devices, based upon data acquired when the given device is inactive; and
comparing the first and the second regression models to determine whether the given
device is relevant to the monitored resource.

11. The method of Claim 10, wherein said comparing step comprises the step of
determining whether the first and the second regression models are statistically equivalent for a
same number of active devices other than the given device.

12. The method of Claim 1, wherein said predicting step comprises the steps of:
dividing the plurality of devices into device classes; and
counting the number of active devices in each of the device classes.

13. The method of Claim 12, further comprising the step of fitting a prediction model for a monitored resource, wherein prediction model parameters of the prediction model depend on the number of active devices in each of the device classes.

14. The method of Claim 12, further comprising the step of computing a prediction model for the number of active devices in each of the device classes.

15. The method of Claim 1, wherein said predicting step comprises the steps of:
computing a regression model of prediction parameters on the number of active devices;
constructing a prediction model for a distribution of the number of active devices; and
combining the regression model and the prediction model for the distribution of active devices to produce a prediction model for one or more of the prediction parameters.

16. The method of Claim 15, wherein said predicting step further comprises the steps of:
computing a regression model of prediction parameters on the number of active devices;
constructing an empirical distribution of the number of active devices; and
combining the regression model and the empirical distribution to produce a prediction model.

17. The method of Claim 16, wherein said step of combining the regression model and the empirical distribution comprises the step of computing an expected value of the prediction parameters with respect to the empirical distribution.

18. The method of Claim 17, wherein said combining step further comprises the steps of:

for each of the one or more prediction parameters,

for each of the monitored values of the number of active devices, computing confidence intervals for the one or more prediction parameters; and

selecting a corresponding one of the confidence intervals for the expected value of each of the one or more prediction parameters.

19. The method of Claim 17, wherein said combining step further comprises the steps of:

computing confidence intervals for the prediction parameters for each of the monitored values of the number of active devices; and

combining the confidence intervals for the prediction parameters for each of the monitored values of the number of active devices with the empirical distribution of the number of active devices, to produce modified confidence intervals for the prediction parameters.

20. The method according to Claim 1, wherein said method is implemented by a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform said method steps.

21. A method for identifying resource saturation in a computer system having a plurality of devices, comprising the steps of:

monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.

22. The method of Claim 21, wherein said identifying step comprises the steps of: fitting a regression model of a monitored resource on the number of active devices; and detecting, in the regression model, departures of the monitored resource from linearity.

23. The method of Claim 21, wherein said identifying step comprises the steps of: fitting a regression model of prediction parameters of the monitored resource on the number of active devices; and detecting departures from linearity of the prediction parameters of the monitored resource.

24. The method of Claim 21, wherein said identifying step comprises the steps of: identifying any of the plurality of devices that are relevant to a monitored resource; and restricting at least one subsequent operation of the computer system that corresponds to the monitored resource to use only devices identified as relevant to the monitored resource from among the plurality of devices.

25. The method of Claim 24, wherein said identifying step comprises the steps of: for a given device currently being evaluated for relevance to the monitored resource,

computing a first regression model for the monitored resource on the number of active devices, based upon data acquired when the given device is active;

computing a second regression model for the monitored resource on the number of active devices, based upon data acquired when the given device is inactive; and

comparing the first and the second regression models to determine whether the given device is relevant to the monitored resource.

26. The method of Claim 24, wherein said comparing step comprises the step of determining whether the first and the second regression models are statistically equivalent for a same number of active devices other than the given device.

27. The method of Claim 21, wherein said identifying step comprises the steps of:
dividing the plurality of devices into device classes; and
counting the number of active devices in each of the device classes.

28. The method of Claim 27, further comprising the step of fitting a prediction model for a monitored resource, wherein prediction model parameters of the prediction model depend on the number of active devices in each of the device classes.

29. The method of Claim 27, further comprising the step of computing a prediction model for the number of active devices in each of the device classes.

30. The method according to Claim 21, wherein said method is implemented by a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform said method steps.

31. A method for predicting effects of adding a new device on a computer system having a plurality of devices, the method comprising the steps of:
monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and
predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices.

32. The method of Claim 31, wherein said predicting step comprises the steps of:
selecting a monitored resource; and
predicting the effects of adding the new device with respect to the selected monitored resource, based upon the monitored values of the resource utilization and the number of active devices.

33. The method of Claim 32, wherein said predicting step comprises the steps of:
constructing a first prediction model of a distribution of the number of active devices; and
modifying the first prediction model to produce a modified prediction model of the distribution of the number of active devices that accounts for the new device.

34. The method of Claim 33, further comprising the steps of:

computing a first prediction model for the selected monitored resource, based upon the first prediction model of the distribution of the number of active devices;

producing a first prediction of the selected monitored resource using the first prediction model for the selected monitored resource;

computing a modified prediction model for the selected monitored resource to account for the new device, based upon the modified prediction model of the distribution of the number of active devices;

producing a second prediction of the selected monitored resource, based upon the modified prediction model for the selected monitored resource; and

comparing the first prediction and the second prediction of the selected monitored resource to evaluate the effects of adding the new device.

35. The method of Claim 31, wherein said predicting step comprises the step of:
dividing the plurality of devices into device classes; and
counting the number of active devices in each of the device classes.

36. The method of Claim 32, further comprising the steps of:
identifying any of the plurality of devices that are relevant to a monitored resource; and
restricting at least one subsequent operation of the computer system that corresponds to the monitored resource to use only devices identified as relevant to the monitored resource from among the plurality of devices.

37. The method according to Claim 31, wherein said method is implemented by a program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform said method steps.

38. An apparatus for predicting a subsequent resource utilization in a computer system having a plurality of devices, comprising:

a monitoring device for monitoring, over a period of time, a contemporaneous resource utilization and a number of active devices to obtain monitored values of the contemporaneous resource utilization and the number of active devices; and

a prediction device for predicting the subsequent resource utilization, based upon the monitored values of the contemporaneous resource utilization and the number of active devices.

39. An apparatus for identifying resource saturation in a computer system having a plurality of devices, comprising:

a monitoring device for monitoring, over a period of time, resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

a forecasting device for identifying resource saturation, based upon the monitored values of the resource utilization and the number of active devices.

40. An apparatus for predicting effects of adding a new device on a computer system having a plurality of devices, the apparatus comprising:

a monitoring device for monitoring, over a period of time, a resource utilization and a number of active devices to obtain monitored values of the resource utilization and the number of active devices; and

a forecasting device for predicting the effects of adding the new device, based upon the monitored values of the resource utilization and the number of active devices.

41. (cancelled)

42. A method for predicting a subsequent resource utilization in a computer system having a plurality of devices, the plurality of devices comprising active devices and non-active devices, comprising the steps of:

monitoring, over a period of time, a contemporaneous resource utilization to obtain first monitored values of the contemporaneous resource utilization;

monitoring, over the period of time, a number of the active devices to obtain second monitored values of the number of the active devices, wherein the monitored number is capable of varying over the period of time;

monitoring, over the period of time, a type of each of the active devices to obtain third monitored values of the type of the each of the active devices; and

predicting the subsequent resource utilization, based upon the first monitored values, the second monitored values, and the third monitored values.

EVIDENCE APPENDIX

None

RELATED PROCEEDINGS APPENDIX

None